

```

    y2=(0.5/wi1)*(y(i)-y(n-i+1))
    y(i)=0.5*(y1+y2)
    y(n-i+1)=0.5*(y1-y2)
    wtemp=wr1
    wr1=wr1*wpr-wi1*wpi+wr1
    wi1=wi1*wpr+wtemp*wpi+wi1
  enddo i6
endif
return
END

```

An alternative way of implementing this algorithm is to form an auxiliary function by copying the even elements of f_j into the first $N/2$ locations, and the odd elements into the next $N/2$ elements in reverse order. However, it is not easy to implement the alternative algorithm without a temporary storage array and we prefer the above in-place algorithm.

Finally, we mention that there exist fast cosine transforms for small N that do not rely on an auxiliary function or use an FFT routine. Instead, they carry out the transform directly, often coded in hardware for fixed N of small dimension [1].

CITED REFERENCES AND FURTHER READING:

- Brigham, E.O. 1974, *The Fast Fourier Transform* (Englewood Cliffs, NJ: Prentice-Hall), §10–10.
 Sorensen, H.V., Jones, D.L., Heideman, M.T., and Burris, C.S. 1987, *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. ASSP-35, pp. 849–863.
 Hou, H.S. 1987, *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. ASSP-35, pp. 1455–1461 [see for additional references].
 Hockney, R.W. 1971, in *Methods in Computational Physics*, vol. 9 (New York: Academic Press).
 Temperton, C. 1980, *Journal of Computational Physics*, vol. 34, pp. 314–329.
 Clarke, R.J. 1985, *Transform Coding of Images*, (Reading, MA: Addison-Wesley).
 Gonzalez, R.C., and Wintz, P. 1987, *Digital Image Processing*, (Reading, MA: Addison-Wesley).
 Chen, W., Smith, C.H., and Fralick, S.C. 1977, *IEEE Transactions on Communications*, vol. COM-25, pp. 1004–1009. [1]

12.4 FFT in Two or More Dimensions

Given a complex function $h(k_1, k_2)$ defined over the two-dimensional grid $0 \leq k_1 \leq N_1 - 1$, $0 \leq k_2 \leq N_2 - 1$, we can define its two-dimensional discrete Fourier transform as a complex function $H(n_1, n_2)$, defined over the same grid,

$$H(n_1, n_2) \equiv \sum_{k_2=0}^{N_2-1} \sum_{k_1=0}^{N_1-1} \exp(2\pi i k_2 n_2 / N_2) \exp(2\pi i k_1 n_1 / N_1) h(k_1, k_2) \quad (12.4.1)$$

By pulling the “subscripts 2” exponential outside of the sum over k_1 , or by reversing the order of summation and pulling the “subscripts 1” outside of the sum over k_2 ,

we can see instantly that the two-dimensional FFT can be computed by taking one-dimensional FFTs sequentially on each index of the original function. Symbolically,

$$\begin{aligned} H(n_1, n_2) &= \text{FFT-on-index-1}(\text{FFT-on-index-2}[h(k_1, k_2)]) \\ &= \text{FFT-on-index-2}(\text{FFT-on-index-1}[h(k_1, k_2)]) \end{aligned} \quad (12.4.2)$$

For this to be practical, of course, both N_1 and N_2 should be some efficient length for an FFT, usually a power of 2. Programming a two-dimensional FFT, using (12.4.2) with a one-dimensional FFT routine, is a bit clumsier than it seems at first. Because the one-dimensional routine requires that its input be in consecutive order as a one-dimensional complex array, you find that you are endlessly copying things out of the multidimensional input array and then copying things back into it. This is not recommended technique. Rather, you should use a multidimensional FFT routine, such as the one we give below.

The generalization of (12.4.1) to more than two dimensions, say to L -dimensions, is evidently

$$\begin{aligned} H(n_1, \dots, n_L) &\equiv \sum_{k_L=0}^{N_L-1} \cdots \sum_{k_1=0}^{N_1-1} \exp(2\pi i k_L n_L / N_L) \times \cdots \\ &\quad \times \exp(2\pi i k_1 n_1 / N_1) h(k_1, \dots, k_L) \end{aligned} \quad (12.4.3)$$

where n_1 and k_1 range from 0 to $N_1 - 1, \dots, n_L$ and k_L range from 0 to $N_L - 1$. How many calls to a one-dimensional FFT are in (12.4.3)? Quite a few! For each value of k_1, k_2, \dots, k_{L-1} you FFT to transform the L index. Then for each value of k_1, k_2, \dots, k_{L-2} and n_L you FFT to transform the $L - 1$ index. And so on. It is best to rely on someone else having done the bookkeeping for once and for all.

The inverse transforms of (12.4.1) or (12.4.3) are just what you would expect them to be: Change the i 's in the exponentials to $-i$'s, and put an overall factor of $1/(N_1 \times \cdots \times N_L)$ in front of the whole thing. Most other features of multidimensional FFTs are also analogous to features already discussed in the one-dimensional case:

- Frequencies are arranged in wrap-around order in the transform, but now for each separate dimension.
- The input data are also treated as if they were wrapped around. If they are discontinuous across this periodic identification (in any dimension) then the spectrum will have some excess power at high frequencies because of the discontinuity. The fix, if you care, is to remove multidimensional linear trends.
- If you are doing spatial filtering and are worried about wrap-around effects, then you need to zero-pad all around the border of the multidimensional array. However, be sure to notice how costly zero-padding is in multidimensional transforms. If you use too thick a zero-pad, you are going to waste a *lot* of storage, especially in 3 or more dimensions!
- Aliasing occurs as always if sufficient bandwidth limiting does not exist along one or more of the dimensions of the transform.

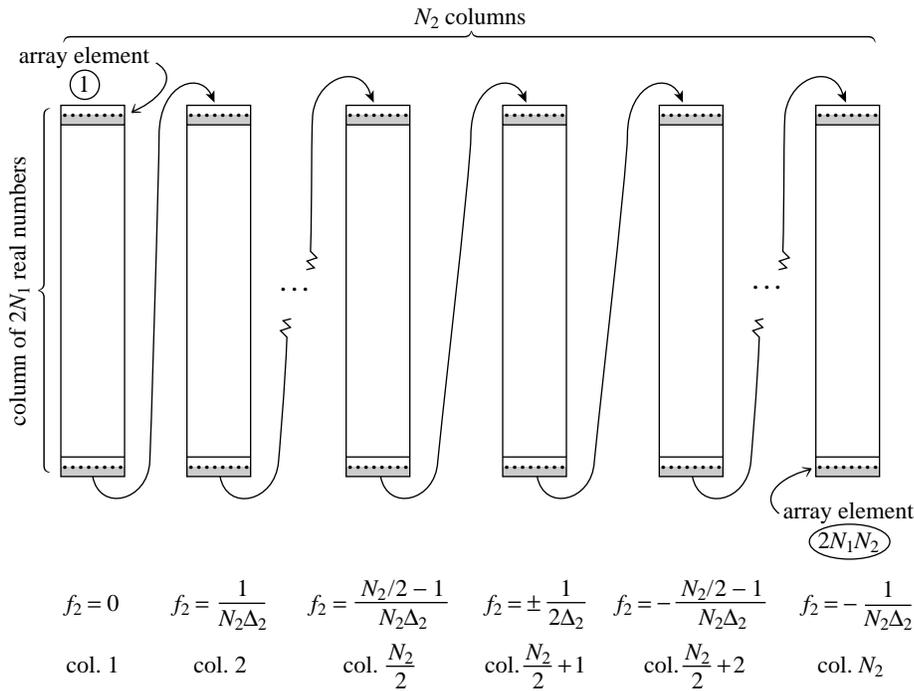


Figure 12.4.1. Storage arrangement of frequencies in the output $H(f_1, f_2)$ of a two-dimensional FFT. The input data is a two-dimensional $N_1 \times N_2$ array $h(t_1, t_2)$ (stored by columns of complex numbers). The output is also stored by complex columns. Each column corresponds to a particular value of f_2 , as shown in the figure. Within each column, the arrangement of frequencies f_1 is exactly as shown in Figure 12.2.2 Δ_1 and Δ_2 are the sampling intervals in the 1 and 2 directions, respectively. The total number of (real) array elements is $2N_1N_2$. The program `fourn` can also do more than two dimensions, and the storage arrangement generalizes in the obvious way.

The routine `fourn` that we furnish herewith is a descendant of one written by N. M. Brenner. It requires as input (i) a scalar, telling the number of dimensions, e.g., 2; (ii) a vector, telling the length of the array in each dimension, e.g., (32,64). Note that these lengths *must all* be powers of 2, and are the numbers of *complex* values in each direction; (iii) the usual scalar equal to ± 1 indicating whether you want the transform or its inverse; and, finally (iv) the array of data.

A few words about the data array: `fourn` accesses it as a one-dimensional array of real numbers, of length equal to twice the product of the lengths of the L dimensions. It assumes that the array represents an L -dimensional complex array, in normal FORTRAN order. Normal FORTRAN order means: (i) each complex value occupies two sequential locations, real part followed by imaginary; (ii) the first subscript changes most rapidly as one goes through the array; the last subscript changes least rapidly; (iii) subscripts range from 1 to their maximum values (N_1, N_2, \dots, N_L , respectively), rather than from 0 to $N_1 - 1, N_2 - 1, \dots, N_L - 1$. Almost all failures to get `fourn` to work result from improper understanding of the above ordering of the data array, so take care! (Figure 12.4.1 illustrates the format of the output array.)

Sample page from NUMERICAL RECIPES IN FORTRAN 77: THE ART OF SCIENTIFIC COMPUTING (ISBN 0-521-43064-X)
 Copyright (C) 1986-1992 by Cambridge University Press. Programs Copyright (C) 1986-1992 by Numerical Recipes Software.
 Permission is granted for internet users to make one paper copy for their own personal use. Further reproduction, or any copying of machine-readable files (including this one), to any server computer, is strictly prohibited. To order Numerical Recipes books or CDROMs, visit website <http://www.nr.com> or call 1-800-872-7423 (North America only), or send email to directcustserv@cambridge.org (outside North America).

```

SUBROUTINE foun(data,nn,ndim,isign)
INTEGER isign,ndim,nn(ndim)
REAL data(*)
  Replaces data by its ndim-dimensional discrete Fourier transform, if isign is input as
  1. nn(1:ndim) is an integer array containing the lengths of each dimension (number of
  complex values), which MUST all be powers of 2. data is a real array of length twice the
  product of these lengths, in which the data are stored as in a multidimensional complex
  FORTRAN array. If isign is input as -1, data is replaced by its inverse transform times
  the product of the lengths of all dimensions.
INTEGER i1,i2,i2rev,i3,i3rev,ibit,idim,ifp1,ifp2,ip1,ip2,
*   ip3,k1,k2,n,nprev,nrem,ntot
REAL tempi,tempr
DOUBLE PRECISION theta,wi,wpi,wpr,wr,wtemp      Double precision for trigonometric re-
ntot=1                                           currences.
do 11 idim=1,ndim                               Compute total number of complex values.
  ntot=ntot*nn(idim)
enddo 11
nprev=1
do 18 idim=1,ndim                               Main loop over the dimensions.
  n=nn(idim)
  nrem=ntot/(n*nprev)
  ip1=2*nprev
  ip2=ip1*n
  ip3=ip2*nrem
  i2rev=1
do 14 i2=1,ip2,ip1                             This is the bit-reversal section of the routine.
  if(i2.lt.i2rev)then
    do 13 i1=i2,i2+ip1-2,2
      do 12 i3=i1,ip3,ip2
        i3rev=i2rev+i3-i2
        tempr=data(i3)
        tempi=data(i3+1)
        data(i3)=data(i3rev)
        data(i3+1)=data(i3rev+1)
        data(i3rev)=tempr
        data(i3rev+1)=tempi
      enddo 12
    enddo 13
  endif
  ibit=ip2/2
1  if ((ibit.ge.ip1).and.(i2rev.gt.ibit)) then
    i2rev=i2rev-ibit
    ibit=ibit/2
    goto 1
  endif
  i2rev=i2rev+ibit
enddo 14
ifp1=ip1                                       Here begins the Danielson-Lanczos section of the routine.
2  if(ifp1.lt.ip2)then
    ifp2=2*ifp1
    theta=isign*6.28318530717959d0/(ifp2/ip1)   Initialize for the trig. recur-
    wpr=-2.d0*sin(0.5d0*theta)**2              rence.
    wpi=sin(theta)
    wr=1.d0
    wi=0.d0
do 17 i3=1,ifp1,ip1
  do 16 i1=i3,i3+ip1-2,2
    do 15 i2=i1,ip3,ifp2
      k1=i2   Danielson-Lanczos formula:
      k2=k1+ifp1
      tempr=sngl(wr)*data(k2)-sngl(wi)*data(k2+1)
      tempi=sngl(wr)*data(k2+1)+sngl(wi)*data(k2)
      data(k2)=data(k1)-tempr
      data(k2+1)=data(k1+1)-tempi
    enddo 15
  enddo 16
enddo 17

```

Sample page from NUMERICAL RECIPES IN FORTRAN 77: THE ART OF SCIENTIFIC COMPUTING (ISBN 0-521-43064-X)
 Copyright (C) 1986-1992 by Cambridge University Press. Programs Copyright (C) 1986-1992 by Numerical Recipes Software.
 Permission is granted for internet users to make one paper copy for their own personal use. Further reproduction, or any copying of machine-
 readable files (including this one), to any server computer, is strictly prohibited. To order Numerical Recipes books or CDROMs, visit website
<http://www.nr.com> or call 1-800-872-7423 (North America only), or send email to directcustserv@cambridge.org (outside North America).

```

        data(k1)=data(k1)+tempr
        data(k1+1)=data(k1+1)+tempi
    enddo 15
enddo 16
wtemp=wr          Trigonometric recurrence.
wr=wr*wpr-wi*wpi+wr
wi=wi*wpr+wtemp*wpi+wi
enddo 17
ifp1=ifp2
goto 2
endif
nprev=n*nprev
enddo 18
return
END

```

CITED REFERENCES AND FURTHER READING:

Nussbaumer, H.J. 1982, *Fast Fourier Transform and Convolution Algorithms* (New York: Springer-Verlag).

12.5 Fourier Transforms of Real Data in Two and Three Dimensions

Two-dimensional FFTs are particularly important in the field of image processing. An image is usually represented as a two-dimensional array of pixel intensities, real (and usually positive) numbers. One commonly desires to filter high, or low, frequency spatial components from an image; or to convolve or deconvolve the image with some instrumental point spread function. Use of the FFT is by far the most efficient technique.

In three dimensions, a common use of the FFT is to solve Poisson's equation for a potential (e.g., electromagnetic or gravitational) on a three-dimensional lattice that represents the discretization of three-dimensional space. Here the source terms (mass or charge distribution) and the desired potentials are also real. In two and three dimensions, with large arrays, memory is often at a premium. It is therefore important to perform the FFTs, insofar as possible, on the data "in place." We want a routine with functionality similar to the multidimensional FFT routine `fourn` (§12.4), but which operates on real, not complex, input data. We give such a routine in this section. The development is analogous to that of §12.3 leading to the one-dimensional routine `realft`. (You might wish to review that material at this point, particularly equation 12.3.5.)

It is convenient to think of the independent variables n_1, \dots, n_L in equation (12.4.3) as representing an L -dimensional vector \vec{n} in wave-number space, with values on the lattice of integers. The transform $H(n_1, \dots, n_L)$ is then denoted $H(\vec{n})$.

It is easy to see that the transform $H(\vec{n})$ is periodic in each of its L dimensions. Specifically, if $\vec{P}_1, \vec{P}_2, \vec{P}_3, \dots$ denote the vectors $(N_1, 0, 0, \dots)$, $(0, N_2, 0, \dots)$, $(0, 0, N_3, \dots)$, and so forth, then

$$H(\vec{n} \pm \vec{P}_j) = H(\vec{n}) \quad j = 1, \dots, L \quad (12.5.1)$$

Sample page from NUMERICAL RECIPES IN FORTRAN 77: THE ART OF SCIENTIFIC COMPUTING (ISBN 0-521-43064-X)
 Copyright (C) 1986-1992 by Cambridge University Press. Programs Copyright (C) 1986-1992 by Numerical Recipes Software.
 Permission is granted for internet users to make one paper copy for their own personal use. Further reproduction, or any copying of machine-readable files (including this one), to any server computer, is strictly prohibited. To order Numerical Recipes books or CDROMs, visit website <http://www.nr.com> or call 1-800-872-7423 (North America only), or send email to directcustserv@cambridge.org (outside North America).